

# Welcome to the HEC-IWG File Systems and I/O Research Guidance Workshop

08/16/05

Grapevine Texas

Gary Grider of Los Alamos  
for the HEC-IWG

ADC-Reviewed

Safety and Security first!

Know your exits!  
Be safe!

Do not tell us anything that cant be  
globally shared!

# What are the goals of this workshop?

- Provide the first input into a multi-year, multi-agency effort to coordinate government spending on research into the HEC-IWG focus area, I/O and File Systems
  - Catalog existing relevant research in this area (high level only) funded by the government, industry, and higher education
  - Look for potential overlaps and gaps
  - List the top areas that need to be addressed in the coming years both short and long term
  - Recommend
    - Areas that are receiving too much emphasis or that will be sufficiently covered without government assistance
    - Most pressing future short term research needs
    - Best bets for long term research needs
    - Any actions needed to stimulate the next new big ideas in this critical area
    - As soon as possible, a list of short and long term research topics to possibly be included in a near future National Science Foundation (NSF) call for research

# What are NOT the goals of this workshop?

- This is NOT an
  - archiving workshop,
  - nor is it a high level data management workshop,
  - nor is it a Grid workshop,
  - however it will touch on those topics as they relate to file systems and middleware.
- This is NOT a platform to talk about specific current products or solutions
- This is NOT about marketing products or solutions
- This is NOT about marketing your institutions research capabilities
- This is NOT a bragging session to see who has the most bytes or bytes/sec
- This is NOT about about current product needs, what you would like to see a product implement, etc.
- This is NOT a place to look for development dollars

# Who was invited and why and roles?

- Agencies/Sites
  - HEC related agencies and sites were invited to share
    - Current research efforts being funded
    - Perceived needed future research efforts based on a 5-10 year view of your sites requirements
    - NOT current deployment, development, or management, efforts
    - NOT why you like a product, don't like a product, how you want a product to change

# Who was invited and why and roles?

- Top research institutions in this area
  - Universities with major emphasis or center of excellence in the file systems and I/O area
    - We are interested in
      - Current research in this area
      - What research people are asking you for
      - What research you think people will ask you for
      - What areas of research you plan to gear up for
      - NOT current near term development efforts
      - NOT an advertisement for your institutions capabilities to provide research

# Who was invited and why and roles?

- Select Industry Representatives
  - ONLY companies that **fund** file systems and I/O research, particularly at the top research institutions
    - All companies invited came from list of funders of top research institutions
      - <http://www.pdl.cmu.edu/affiliates.html>, <http://ssrc.cse.ucsc.edu/sponsors.shtml>, <http://www.dtc.umn.edu/industrial/affiliateslist.html> , etc.
  - Why - initially there were no industry participants, but we felt that if there were companies that are clearly acting as peers to the government in funding research, we need this information to know what the government should fund and not fund
    - Several companies that don't fund research institutions wanted to be invited
    - Several companies that do product related research wanted to be invited
  - We really wanted only those companies that are acting as peers to the government that fund generally applicable research
  - We are Interested in research you fund now, what you plan to fund, what areas you are likely to concentrate on from a research funding point of view, etc. We request that you put on your industry funding research hat as much as possible.
  - This is NOT about products, product development, not even GPFS, PanFS, Lustre, QFS, your favorite disk array, your favorite SAN, etc.

# The magic questions!

- 1) Current high level topics of File Systems and I/O Research being done
- 2) Areas that need to have more focus (gaps)
- 3) Areas that have or will have too much focus (overlaps/other)
- 4) What you would recommend the government focus on for funded research in the coming 3-10 years, short term and long term.



# Agenda

# Breakouts